

CLAIMS

What is claimed is:

- 1 1. A method for applying changes in redo records to make a particular
2 resource reflect changes made to the particular resource in volatile
3 memory before a failure, the method comprising the steps of:
4 establishing links that link together a set of redo records that contain
5 changes made to the particular resource; and
6 following the links to apply the changes contained in the set of redo
7 records to cause the particular resource to reflect the changes made
8 to the particular resource in volatile memory before the failure.

- 1 2. The method of Claim 1, wherein the step of establishing links that link
2 together the set of redo records is performed prior to the failure.

- 1 3. The method of Claim 1, wherein the step of following the links to apply
2 the changes contained in the set of redo records is performed in response
3 to a transaction requesting access to the particular resource subsequent to
4 the failure.

- 1 4. A method for applying changes in redo records to make a resource
2 available, wherein the resource is locked by a dead transaction, the method
3 comprising the steps of:
4 (A) identifying a redo record in a block-base redo chain, wherein the redo
5 record contains least recent changes that need to be applied to the
6 resource;

7 (B) applying the least recent changes contained in the redo record to the
8 resource; and
9 (C) repeating steps (A) and (B) until all changes that are contained in redo
10 records in the block-base redo chain have been applied to the
11 resource.

- 1 5. A method for applying changes in two or more undo records in parallel,
- 2 wherein a plurality of resources are locked by a dead transaction, the
- 3 method comprising the steps of:
- 4 identifying, within a single undo log file, a plurality of sets of undo records,
- 5 wherein each set of undo records of said plurality of sets of undo
- 6 records does not contain any undo record that depends on any undo
- 7 record in any other set of undo records of said plurality of sets of
- 8 undo records; and
- 9 applying the plurality of sets of undo records in parallel relative to one
- 10 another.

1 6. The method of Claim 5, further comprising establishing the plurality of sets
2 of undo records by performing steps that comprise:
3 among the plurality of sets of undo records:
4 assigning all undo records that are associated with a first
5 resource to a first set of undo records; and
6 assigning all undo records that are associated with a
7 second resource to a second set of undo records.

1 7. The method of Claim 6, wherein:

2 the step of assigning all undo records that are associated with the first
3 resource includes the step of linking in a first chain undo records
4 that contain changes that are associated with the first resource; and
5 the step of assigning all undo records that are associated with the second
6 resource includes the step of linking in a second chain undo records
7 that contain changes that are associated with the second resource.

1 8. The method of Claim 7, wherein:
2 the step of linking in the first chain undo records that contain changes that
3 are associated with the first resource includes the step of generating
4 a first block-based undo chain, wherein the first block-based undo
5 chain contains undo records that contain changes that need to be
6 applied to the first resource; and
7 the step of linking in the second chain undo records that contain changes
8 that are associated with the second resource includes the step of
9 generating a second block-based undo chain, wherein the second
10 block-based undo chain contains undo records that contain changes
11 that need to be applied to the second resource.

1 9. A method for linking undo records, the method comprising the steps of:
2 identifying an undo record, wherein the undo record contains change
3 information that is associated with a particular resource;
4 linking the undo record into an undo record chain, wherein the undo record
5 chain contains only undo records that contain change information
6 that is associated with the particular resource; and

7 wherein the step of linking the undo record includes generating identifying
8 data in at least one of the records in the undo record chain or in the
9 undo record, wherein said identifying data once generated identifies
10 a particular record in the undo record chain.

1 10. The method of claim 9, wherein:
2 the step of identifying the undo record includes the step of identifying an
3 undo record that contains change information that is associated with
4 a particular data block; and
5 the step of linking the undo record into the undo record chain, includes the
6 step of linking the undo record into the undo record chain, wherein
7 the undo record chain contains only undo records that contain
8 change information that is associated with the particular data block.

1 11. The method of claim 9, further comprising the steps of:
2 identifying a first undo record, wherein relative to the undo record chain the
3 first undo record contains the least recent change information that
4 needs to be applied to the particular resource; and
5 linking a pointer in the first undo record to the undo record chain.

1 12. The method of claim 9, further comprising the steps of:
2 identifying a last undo record, wherein relative to the undo record chain the
3 last undo record contains the most recent change information that
4 needs to be applied to the particular resource; and
5 linking a pointer in the last undo record to the undo record chain.

1 13. The method of claim 9, wherein the undo record chain contains only undo
2 records that contain change information that needs to be applied to the
3 particular resource.

1 14. A method for recovering after failure of a transaction, the method
2 comprising the step of:
3 prior to said failure, storing in a first recovery record data that reflects a first
4 change made by the transaction to a first resource;
5 after making said first change and prior to said failure, storing in a second
6 recovery record data that reflects a second change by the transaction
7 to a second resource;
8 after making said second change and prior to said failure, storing in a third
9 recovery record data that reflects a third change to said first
10 resource;
11 after said failure, recovering said transaction; and
12 wherein the step of recovering said transaction includes applying said data
13 in said first recovery record and said data in said third recovery
14 record prior to applying said data in said second recovery record.

- 1 15. The method of claim 14, wherein:
 - 2 the step of storing in a first recovery record comprises the step of storing in
 - 3 a first undo record;
 - 4 the step of storing in a second recovery record comprises the step of storing
 - 5 in a second undo record;

6 the step of storing in a third recovery record comprises the step of storing in
7 a third undo record; and
8 the step of applying said data comprises the step of applying said
9 data in said first undo record and said data in said third
10 undo record prior to applying said data in said second
11 undo record.

1 16. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 1.

1 17. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 2.

1 18. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 3.

1 19. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 4.

- 1 20. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 5.

- 1 21. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 6.

- 1 22. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 7.

- 1 23. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 8.

- 1 24. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 9.

- 1 25. A computer-readable medium carrying one or more sequences of
- 2 instructions which, when executed by one or more processors,
- 3 causes the one or more processors to perform the method recited
- 4 in Claim 10.

- 1 26. A computer-readable medium carrying one or more sequences of
- 2 instructions which, when executed by one or more processors,
- 3 causes the one or more processors to perform the method recited
- 4 in Claim 11.

- 1 27. A computer-readable medium carrying one or more sequences of
- 2 instructions which, when executed by one or more processors,
- 3 causes the one or more processors to perform the method recited
- 4 in Claim 12.

- 1 28. A computer-readable medium carrying one or more sequences of
- 2 instructions which, when executed by one or more processors,
- 3 causes the one or more processors to perform the method recited
- 4 in Claim 13.

- 1 29. A computer-readable medium carrying one or more sequences of
- 2 instructions which, when executed by one or more processors,
- 3 causes the one or more processors to perform the method recited
- 4 in Claim 14.

1 30. A computer-readable medium carrying one or more sequences of
2 instructions which, when executed by one or more processors,
3 causes the one or more processors to perform the method recited
4 in Claim 15.